Energy Management Systems and AI: A Review of the Literature

By
Chinweike Ezeokafor
26th August 2021

Energy Management Systems have long been the focus of extensive research. This has led to the development of various techniques for modelling the energy demand during a timeframe. The realisation of the need to reduce energy wastage and costs have contributed to the urgency of this research to be carried out. Moreover, studies from the International Energy Agency have shown that between 2017 and 2022, renewable energy will constitute about 43% of the new net world electricity energy addition [1]. Modelling the energy behaviour of homes is particularly difficult due to the wide range of factors that could affect it on a given day and the fact that if the model underestimates the usage, it could lead to failures in the grid at peak load. This has increased the need for accurate modelling techniques which could predict the energy requirement in an area while also factoring the scalability of the model to cover an increased area. This literary effort will focus on the various techniques and objectives used in the energy management systems in the available literature reviewed. Although the existing body of literature go into detail on other related concepts, this paper will mostly focus on the basic architecture, software used, and the objective of the research reviewed.

The Literature Review

(Han et al.,2014)

This considers the monitoring of the energy consumption in a home using ZigBee based energy management modules. This periodically sends the energy reading from each appliance to which it is connected to a home server. A PLC-based gateway is then used to monitor the energy generated in the home through various sources such as solar panels and other renewable energy sources in the home. This also considers the weather patterns and climatic conditions in an area and its effect on the renewable energy generation in the home. This data is then analysed by the Energy Generation Manager to detect trends in the home energy usage as the weather varies. The home server then modifies the schedule of the appliances in the home to ensure that the bulk of the energy consumption in the home occurs in the high renewable energy generation and low-price time. This enables significant savings to be made on cost and energy usage.The set up can be seen in the figure below

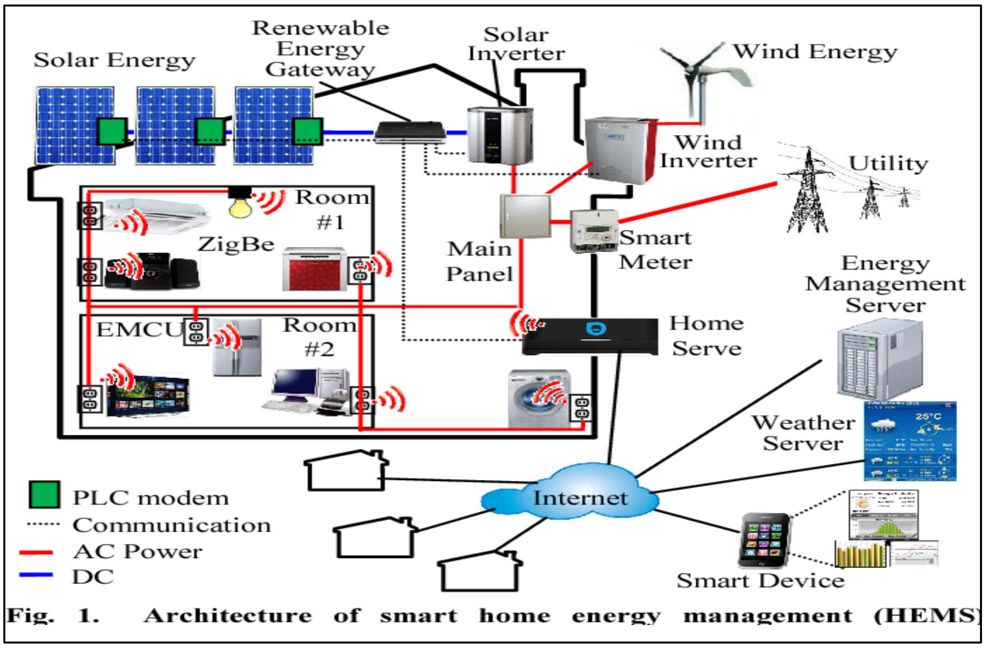

Fig. 1. Architecture of smart home energy management (HEMS)

Figure 1: Aechitecture of Smart Home Energy Management System (HEMS)

The usage patterns of different homes in the area can be analysed and compared against a benchmark to facilitate the home usage pattern analysis . This research is based on the assumption that people would be willing to share their energy usage as this would be required for the comparisons to be made. Further research investigations are necessary to determine how willing users are to share the usage profile in order for the method described above to work. Besides, this research only focuses on sources of fluctuation in the generated electricity and does not consider factors that could affect the energy consumption in a home. The research mainly focuses on the architecture used but does not elaborate on what modelling methods were used to predict the energy consumption and their accuracy [2].

(A.R et al.,2017)

This research direction focuses on the application of energy management systems on the operation of HVAC (Heating, Ventilation and Air Conditioning) as this accounted for 60% of the energy consumption in the home in the region in which the research was carried out. This is done using IoT, Big Data and BI analysis methods. Each of the home appliances is fitted with a module that measures the device's energy consumption and sends it to a centralised server where it is analysed by an off-the-shelf Business Analytic package which then makes decisions based on the data received. The microcontrollers are also fitted with a solid-state relay which allows the devices to be remotely switched off or on by the user or the centralised server. This also displays the energy usage patterns of the home to the user through their smartphones and other platforms. This enables the energy usage within the home to be adjusted, therefore saving both cost and energy. The scalability and security of the energy management system were also discussed, and how it could be implemented, such as hosting the central server on the cloud. This research does not consider what effect, if any, the continuous switching off the appliances would have on their lifespan, especially on devices such as boilers which may not have been designed for continual switching on and off. The limitation of this research also is that it does not consider the sources of energy generation in the home, such as photovoltaics, and its effect on the grid energy consumption [3].

(Tianshu et al.,2018)

(Tianshu et al.,2018), in his work, makes use of Model Predictive Control (MPC) in the automation of energy management systems in buildings. This was used in the co-scheduling of HVAC control, EV charging and battery usage in a home. This method involves building an accurate model of the physical dynamics which could affect the energy consumption in the building. The effectiveness of this method depends on the use of an accurate model which correctly identifies each of the factors affecting the energy consumption and their scope. It was observed that EV charging coincided with peak HVAC use, and this usually occurred during peak load. An MPC based co-scheduling algorithm was used in the peak shaving of the energy demand by load shifting to times with less demand. This was done by leveraging the battery storage and other renewable energy sources while scheduling HVAC and EV charging away from peak time. This was tested in a building with 60 EV charges in a day and achieved a 7.4% reduction in the total energy cost and a 25.4% reduction in peak demand compared to the baseline. One major limitation with the MPC method is the need for accurate models, as this could be both times consuming and expensive to build. Additionally, it might not be possible to consider all the factors that could affect the energy consumption, such as insulation degradation over time [5].

(D. Arcos-Aviles et al.,2018)

This research focuses on smoothing the power profile exchange of a residential grid-connected microgrid (MG) with the grid. This is conducted on a home with renewable energy sources such as wind and photovoltaic and batteries for energy storage. The loads from the house were fixed and uncontrollable, and it was assumed that only the battery State of Charge (SoC) could be controlled. The main objective of this research was the smoothing the power exchanged with the grid while concurrently satisfying at any time the load demand (i.e., there was no demand-side management) and the ESS constraints. The optimisation and smoothing were done using a low-complexity Fuzzy Logic Controller design with a two-input one-output Fuzzy Controller with 25 rules. Fuzzy logic was used because it allowed the system to analyse different inputs and decide based on how near the inputs were to the optimal or desired value. This was found to be an easy way to incorporate renewable energy into a conventional installation by simply adding a set of batteries and a renewable source [6].

(Lissa et al., 2021)

In this research, (Lissa et al., 2021) focuses more on finding a more efficient machine-learning algorithm to enable the optimisation of photovoltaic (PV) energy consumption of smart homes. These homes are assumed to have automated actuators and sensors for monitoring already and controlling the energy consumption of the appliances in the home. The algorithm used here is the Deep Reinforming Learning algorithm. This algorithm was chosen because it does not require prior information about the building to optimise the energy consumption. This had a significant benefit over other methods such as fuzzy control or Model Predictive Control (MPC), which relies on building an accurate model, which could be time-consuming and error-prone. It was shown to achieve an average of 8% saving against the baseline consumption. This was seen to rise to 16% in the summer months. The optimisation process involved balancing user comfort against energy savings and was seen to balance this with more emphasis given on user comfort effectively. It also resulted in a 9.5% increase in the utilisation of renewable energy compared to a rule-based approach. In addition, by effectively predicting when loads are added, it was able to load shift to periods that are outside peak usage time, reducing the peak [4]

Conclusion

This research focuses on finding better techniques for running Home Energy Management System. The purpose of this review is to analyse the various methods and techniques used and their development over time. One area of further analysis is the perception of the public to such technology with emphasis on the security of the vast amount of data gathered and its impact on the privacy of the users. Moreover, as most of the work conducted in this research focused on advanced nations with stable grids, further research could be done to determine the best technique for energy management systems in developing nations or secluded rural areas.

References

[1]. (a). International Energy Agency (IEA). Renewables 2017: Analysis and Forecasts to 2022, 2017. https://doi.org/10.1787/re_mar-2017-en.

[2]. Han, C. Choi, W. Park, I. Lee and S. Kim, "Smart home energy management system including renewable energy based on ZigBee and PLC," in *IEEE Transactions on Consumer Electronics*, vol. 60, no. 2, pp. 198-202, May 2014, doi: 10.1109/TCE.2014.6851994.

[3]. A. R. Al-Ali, I. A. Zualkernan, M. Rashid, R. Gupta and M. Alikarar, "A smart home energy management system using IoT and big data analytics approach," in IEEE Transactions on Consumer Electronics, vol. 63, no. 4, pp. 426-434, November 2017, doi: 10.1109/TCE.2017.015014

[4]. Lissa, P. et al., 2021. Deep reinforcement learning for home energy management system control. Energy and AI, 3, p.100043

[5]. Tianshu Wei, Xiaoming Chen, Xin Li, and Qi Zhu. 2018. Model-based and data-driven approaches for building automation and control. In Proceedings of the International Conference on Computer-Aided Design (ICCAD '18). Association for Computing Machinery, New York, NY, USA, Article 26, 1–8. DOI: https://doi.org/10.1145/3240765.3243485

[6]. D. Arcos-Aviles, J. Pascual, L. Marroyo, P. Sanchis and F. Guinjoan, "Fuzzy Logic-Based Energy Management System Design for Residential Grid-Connected Microgrids," in IEEE Transactions on Smart Grid, vol. 9, no. 2, pp. 530-543, March 2018, doi: 10.1109/TSG.2016.2555245.

Deep Reinforcement Learning Algorithm: A Review of the Literature

By

Chinweike Ezeokafor

26th August 2021

After an extensive review of energy management systems and the techniques employed to manage them, Deep Reinforcement Learning Algorithm was seen as the most promising technique for energy management systems. This review begins with an introduction to the concept of Deep Reinforcement Learning Algorithm. Reinforcement Learning is deeply rooted in the psychological and neuroscientific perspective of animal behaviour and how they can optimise their behaviour to get the best result in each environment [3]. Reinforcement learning builds upon this by using an artificial agent. This agent learns by interacting with its environment, and then through the realised experience , optimises an objective which is given in the form of cumulative rewards. The agent learns about its environment through exploration and then exploit those decisions that were seen to lead it to its objective. At the start, the agent does not know much about its environment, therefore explores it. As it gets more confident, it interacts with the environment through exploitation. The balance must be reached to ensure the agent can sufficiently explore its environment to discover more rewarding actions (better solutions) while leveraging on those proven to work in the past. The environment might be stochastic (the agent is not given complete information about the environment), or it may be high-dimensional. Deep Learning involves the use of computational models that are composed of multiple processing layers. These methods allow the representation of data with multiple levels of abstraction to be learnt[2]. These methods have enabled great leaps in speech recognition, autonomous vehicles, visual image processing, object detection and genomics due to their ability to discover intricate structures in high dimensional data, therefore enabling significant advancement in problems that have plagued the artificial intelligence community for years. Deep Reinforcement Learning Algorithm involves the use of Deep Neural Networks and Reinforcement Learning to solve challenging computational and optimisation tasks which have been unsolvable prior , using other machine learning techniques [1]. It has also enabled the reduction of the effect of the curse of dimensionality. This is where the number of calculations required for a task increases drastically with the increase in input, as demonstrated by Minh et al. in 2015 [12].

This review focuses on the various applications of the Deep Reinforcement Learning Algorithm, the techniques employed and the challenges facing this method. This is done considering its usefulness as a means of optimising an energy management system. This review covers the progress achieved in the field of machine learning and artificial intelligence in the field of Deep Reinforcement Learning and how this can be applied in the optimisation of energy management systems in an intelligent home

Applications of Deep Neural Networks and Reinforcement Learning

This review evaluates some of the research conducted in the field of Deep Reinforcement Learning, the applications and results outcome.

Language Translation

(Bengio et al., 2008) proposes the use of neural-based statistical language models to deal with language translation. Statistical language techniques have been limited in their application due to the high dimensionality of each language which can contain up to hundreds of thousands of words in their vocabulary and language sparsity. This research proposes a method to increase the speed at which the model is trained using adaptive importance sampling. This is based on the observation that the gradient of the log-likelihood can be separated into two parts: positive and negative contributions. The negative contributions were then estimated through importance sampling. This was seen to achieve a 19-fold speedup when compared to the standard setup. An even more significant speedup of about 150 was achieved by adapting the proposal distribution as training progressed, so it stayed as close as possible to the network's distribution. One major issue encountered during the research was the fact that it was difficult to make an efficient parallel implementation of the sampling algorithm. This was because parallelisation had to be done in a hidden layer which meant that for each backpropagation, a gradient had to be accumulated with respect to the feature parameters for each processor, which then had to be shared. The challenge arises when the gradient is shared as these were seen to necessitate enormous resources and was found to take up to 60% of backpropagation time. This was reduced through desynchronising the sharing of the parameters on the feature vectors. Although this affected convergence, the effect was seen to be insignificant. Another problem that was faced was the choosing of an Effective Sample Size (ESS) that would guarantee convergence. (Wu et al., 2018) proposes a means of training Neural Machine Translation (NMT) models more efficiently through the use of Reinforcement Learning. This was seen to achieve a BLEU score of 26.73, even surpassing the best ensemble model in WMT17 Chinese-English translation challenges. The instability in gradient estimation and reward computation was reduced through using a function learning approach and simple networks to build the learning function. An area suggested for further research include looking at the factors that affect the performance of the proposed adaptive importance sampling method.

Optimisation of Gameplay-Atari 2600

(Mnih et al.,2013) focuses on the use of deep learning and reinforcement learning algorithms to solve complex problems involving high-level sensory data. Some of the challenges faced while using deep neural networks with supervised learning (SL) includes the fact that deep learning requires a large amount of hand-labelled training data while reinforcement learning trains using a scalar reward signal and learns from interactions with its environment. Also, in reinforcement learning, the data distribution changes as the algorithm learn new techniques, and this is problematic for deep learning as it is assumed that the data have a fixed underlying distribution. The researchers overcame this problem through the use of a convolutional neural network that can learn successful control policy from raw data in a complex reinforcement learning environment. An experience replay mechanism was also used to randomly sample the previous iterations, thereby smoothening the training distribution over many past behaviours. This solved the problem of correlated data and non-stationary data distribution in the reinforcement learning environment. This was then applied to a range of Atari 2600 games such as Beam Rider, Breakout, Sea quest and Space Invader and was implemented in The Arcade Learning Environment. The agent was presented with a high dimensional input of 210 x 160 RGB at 60Hz, as well as a diverse set of tasks that were designed to be challenging for a human player. The agent was seen to surpass previous RL techniques on six of the games and even surpassed an expert human player in three of them. These results were achieved without any modification or adjustment of the hyperparameters or architecture as it was used in the different games. One challenge that remains is the evaluation of the critical point from which convergence in the predicted Q-value or other parameters stops and divergence occurs. Although gradient temporal difference methods were used to help prevent this from happening, it does not eradicate it.

Optimisation of Gameplay-AlphaGo

(Silver et al.,2016) investigated the use of deep neural networks and reinforcement learning in innovative ways to solve increasingly challenging problems in their research. This was tested on AlphaGo, which at the time was considered to be one of the most challenging of the classic games for artificial intelligence to solve. This was due to its enormous search space and the difficulties in evaluating board positions and moves in such a space. This was overcome using a clever method that involved training the neural network through a pipeline consisting of several machine learning stages. The previous methods of solving Go mainly involved the use of Monte Carlo tree search (MCTS) enhanced by policies to predict expert human use. However, this limited their maximum performance to those of the trainers and meant that they could never outmatch their human trainers. The first stage involved using supervised learning (SL) policy network directly from expert human use. A fast policy was then trained that rapidly sampled action during rollout. A reinforcement learning policy was then trained and used to improve the network. The improvement was made through using the reinforcement learning policy network to improve the policy gradient of the SL policy network to maximise the outcome (wins) over the previous versions of the policy network. The RL policy network was seen to win about 85% of the game against the SL policy as well as Pachi (previous state of the art algorithm. This was a massive improvement from the previous supervised learning convoluted network, which won 11% of the time against Pachi and 12% against a slightly weaker program, Fuego.

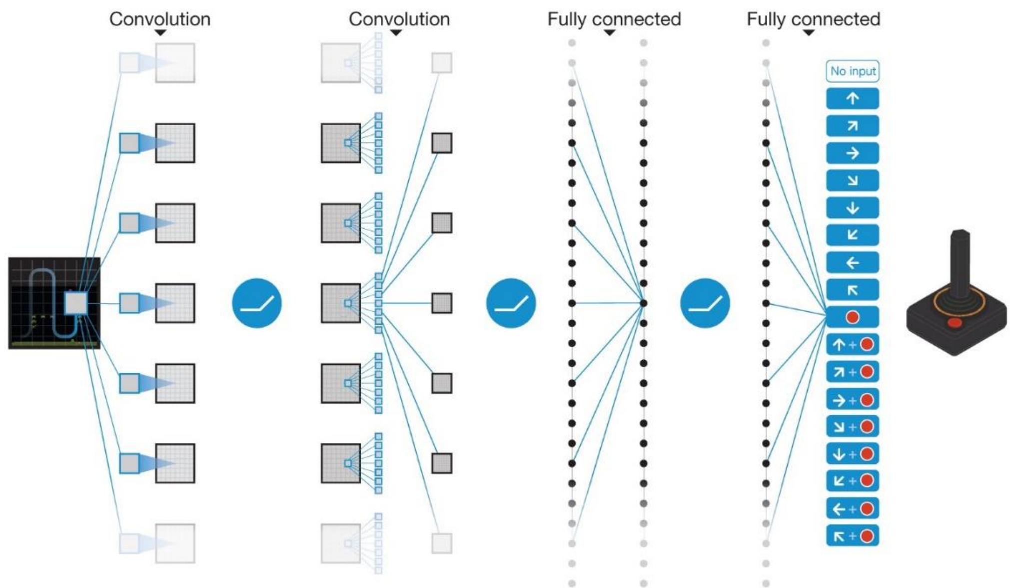

Fig 1: Convoluted Neural Netwok [14]

The program created, called AlphaGo, was then tested in a tournament against other Go programs and was seen to win about 99.8% of the time against other Go programs. The distributed version of AlphaGo was then evaluated against Fan Hui, who was the winner of the 2013, 2014 and 2015 European Go tournament and defeated him in the five games played. A superhuman level performance was achieved using a deep neural network trained by a novel combination of supervised and reinforcement learning.

Control System for Autonomous Vehicles

Faijie et al., 2018 focuses on the use of Deep Reinforcement Learning Algorithms applied with Deep Q Networks to control the actions made by an autonomous vehicle. This research is based on the use of a fully convoluted neural network (CNN) to approximate the Q function. The decisions made by the algorithm is based on input from the front camera image and the LIDAR scanner. Each action made by the algorithm is given a reward, with the most significant positive rewards achieved by achieving greater speed in the positive direction. The most significant negative reward is obtained when the car hits an object. As in RL, present rewards are more valuable than future rewards; the agent is effectively optimised to drive as fast as it can while doing it safely, as any hit leads to a high negative reward. The cumulative reward gained was seen to increase over time, suggesting that it was learning to make the right decisions. (Shalev-Schwartz et al., 2016) focuses on endowing a car with a long-term driving strategy known as "Driving Policy" to enable it to be fully autonomous. The challenge involved endowing robotic autonomous vehicles with human-level negotiation skills while merging, overtaking, and performing other complex actions that rely on visualising and analysing the actions of other road users who could behave in an unexpected manner. Another challenge that was presented in this research is determining the optimal reward (positive and negative) for each task to ensure accident-free driving

Conclusion

This paper analyses the different usage of deep neural networks and reinforcement learning in the real world. They are seen to excel in tasks that involve a high amount of optimisation. These features would allow it to be implemented in an Energy Management System where a reward function could be used to make reducing the amount of energy used from the grid the objective, hence cost reduction. The environment could be created to set specific constraints such as the user's preference for the home's temperature and the specific times in which different appliances must not be turned off. This then presents an optimisation task in which the algorithm would have to achieve the objective for the reward while keeping within the limits of the constraints. These are like those set in Atari 2600, where the algorithm was seen to achieve superhuman performance in a lot of the games.

One of the challenges in deep learning is the difficulty at predicting which deep learning approach works more than the others and how to determine the optimal structure that would perform a specific task best [7]. Nevertheless, the use of reinforcement learning in Energy Management Systems would enable greater efficiency in energy usage to be achieved.

Acknowledgement

The author would like to thank the Laidlaw Foundation for the sponsorship of this research, as well as the guidance provided by his supervisor Dr M. Shahbazi

Reference

[1]. Marr, B., 2021. *What Is Deep Learning AI? A Simple Guide With 8 Practical Examples*. [online] Forbes. Available at: <https://www.forbes.com/sites/bernardmarr/2018/10/01/what-is-deep-learning-ai-a-simple-guide-with-8-practical-examples/?sh=63c7d8f8d4ba> [Accessed 4 October 2021].

[2]. LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning. *nature*, *521*(7553), pp.436-444.

[3]. François-Lavet, V., Henderson, P., Islam, R., Bellemare, M.G. and Pineau, J., 2018. An introduction to deep reinforcement learning. *arXiv preprint arXiv:1811.12560*.

[4]. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. and Riedmiller, M., 2013. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*

[5]. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. and Dieleman, S., 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, *529*(7587), pp.484-489.

[6]. Bengio, Y. and Senécal, J.S., 2008. Adaptive importance sampling to accelerate training of a neural probabilistic language model. *IEEE Transactions on Neural Networks*, *19*(4), pp.713-722.

[7]. Guo, Y., Liu, Y., Oerlemans, A., Lao, S., Wu, S. and Lew, M.S., 2016. Deep learning for visual understanding: A review. *Neurocomputing*, *187*, pp.27-48

[8]. A. R. Fayjie, S. Hossain, D. Oualid and D. Lee, "Driverless Car: Autonomous Driving Using Deep Reinforcement Learning in Urban Environment," *2018 15th International Conference on Ubiquitous Robots (UR)*, 2018, pp. 896-901, doi: 10.1109/URAI.2018.844179

[9]. Shalev-Shwartz, S., Shammah, S. and Shashua, A., 2016. Safe, multiagent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*.

[10]. Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction, volume 1. MIT press Cambridge, 2015

[11]. Nguyen, T.T., Nguyen, N.D. and Nahavandi, S., 2020. Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE transactions on cybernetics*, *50*(9), pp.3826-3839

[12]. V. Mnih et al., "Human-level control through deep reinforcement learning", *Nature*, vol. 518, no. 7540, pp. 529-533, 2015.

[13]. Wu, L., Tian, F., Qin, T., Lai, J. and Liu, T.Y., 2018. A study of reinforcement learning for neural machine translation. *arXiv preprint arXiv:1808.08866*.

[14]. Mnih, V., Kavukcuoglu, K., Silver, D. *et al.* Human-level control through deep reinforcement learning. *Nature* **518,** 529–533 (2015). https://doi.org/10.1038/nature14236